

Knowing What I See

ALEX BYRNE

If I descry a hawk, I find the hawk but I do not find my seeing of the hawk. My seeing of the hawk seems to be a queerly transparent sort of process, transparent in that while a hawk is detected, nothing else is detected answering to the verb in 'see a hawk'.

—Ryle, *The Concept of Mind*

I. Introduction

By using my eyes, I can come to know that there is a hawk perching on the fence post. Cognitive science has made enormous progress in understanding how we have this sort of perceptual knowledge. Any textbook on perception will go into detail about the receptors in the eye, the detection of low-level features such as edges, the recovery of 3D shape, and theories of object recognition. Much is controversial, and much is unknown, but the broad shape of a satisfying account, together with some of the intricate small parts, is clear enough.

When I am in a position to know, by using my eyes, that there is a hawk on the fence post, I am usually in a position to know something else, namely that I see a hawk. This second item of knowledge is of course not entailed by the first: idealism aside, the hawk has no essential connection to me or my perceptual state. Equally obviously, this second item of knowledge is not evidentially

probable given the first: the mere fact that there is a hawk on the fence post is hardly good evidence that it is seen, let alone that I am the one who sees it. Which is to say that an account of how I know that there is a hawk on the fence post by vision cannot easily be converted into an account of how I know I see a hawk. Hence, the question this chapter will attempt to answer: How do I know that I see a hawk?¹

Textbooks are of no help, and the issue has received very little discussion by philosophers. The obsession of contemporary epistemology has been my knowledge of the *hawk*, not my knowledge that I *see* a hawk. The explanation of this curious state of affairs is an interesting topic in its own right; it will not be examined here, however.

2. Dretske on Zombies

One philosopher who has addressed our question is Dretske, most extensively in “How Do You Know You Are Not a Zombie?” (2003).² The eponymous question of Dretske’s paper covers perception and awareness in general; it has our question about seeing as a special case. Dretske’s discussion highlights a serious problem with one ostensibly attractive answer to our question, so this is a good place to start.

Dretske’s mention of “zombies” might mislead. In more-or-less standard usage, “zombies” are creatures who are physically exactly like awake and alert human beings, but who are not “phenomenally conscious”—there is “nothing it is like” to be a zombie. Zombies are frequently presumed to have a typical package of intentional mental states. So zombies believe that it is raining, and see hawks, although of course their perceptual states are devoid of any “qualia.” For those who think that this conventional sort of zombie could have existed, the question “How do you know you are not a zombie?” can seem pressing. After all, zombies are (arguably) firmly convinced that they are not zombies—just like us.³

Importantly, Dretskean zombies are not the standard sort, and epistemological issues about qualia are only of peripheral relevance to Dretske’s concerns. In Dretske’s usage, “zombies [are] human-like creatures who are not conscious and,

¹ ‘I see a hawk’ is to be read in the everyday sense in which it entails that there exists a hawk to be seen. If there is a sense of ‘see’ in which I can be truly said to see a phoenix (as is claimed in Anscombe 1965), it does not occur in this chapter.

² See also Shoemaker 1963, especially 83–84.

³ For extensive discussion, see Chalmers 1996, ch. 5. See also note 18 below.

therefore, not conscious of anything” (2003, 9 n. 1). A Dretskean zombie is simply a superficial human look-alike who behaves in humanlike ways and who lacks intentional states; in particular, a Dretskean zombie sees nothing. One day the Sony Corporation will produce mindless robots to help around the house, so sophisticated that the casual observer will take them to be normal humans—Dretskean zombies are rather like that. The possibility of standard zombies is controversial; in contrast, only a hard-line behaviorist would deny that Dretskean zombies could have existed.⁴

2.1. Dretske’s Statement of the Problem

Dretske writes:

In normal (i.e. veridical) perception, then, the objects you are aware of are objective, mind-independent objects. They exist whether or not you experience them. . . . Everything you are aware of would be the same if you were a zombie. In having perceptual experience, then, nothing distinguishes your world, the world you experience, from a zombie’s. This being so, what is it about this world that tells you that, unlike a zombie, you experience it? What is it that you are aware of that indicates that you are aware of it? (2003, 1, note omitted)

Dretske’s point—applied to our running example—was alluded to in this chapter’s second paragraph. The hawk has no special connection to visual experience—provided it does not notice my presence, the bird does not “indicate that I am aware of it.” Hawks can and frequently do happily perch unseen on fence posts; put more generally, the world as revealed by vision does not have vision in it. Thus, the presence of the hawk does not favor the hypothesis that I see it over the “skeptical hypothesis” that I am a (Dretskean) zombie, and hence do not see it. The evidence (facts) provided by vision would be exactly the same even if I were a zombie.⁵

⁴ Dretske briefly alludes to “some readers who doubt that [Dretskean zombies] are possible” (2003, 10 n. 1), so to be on the safe side the explanation of Dretskean zombies in the text should be viewed as a friendly elaboration or amendment.

⁵ Cf. Wittgenstein [1921] 2001, *Tractatus* 5.633, “nothing in the visual field allows you to infer that it is seen by an eye” (referred to in Dretske’s opening paragraph). Wittgenstein’s concern is the self (in fact the “metaphysical subject”), not vision. And the self is another problem: the world as revealed by vision also does not have me in it. (The fact that I can usually see my nose is not the key to this difficulty.) The point is somewhat obscured by concentrating on a Dretskean zombie scenario, which does have me in it (or at any rate a Dretskean zombie who resembles me).

The transparency
proposal

One (vaguely stated) answer to our question is that I know that I see a hawk merely by attending to the scene before my eyes, and in particular to the hawk. Dretske was formerly sympathetic to this idea; with a nod to the epigraph from Ryle, call it the *transparency proposal*.⁶ Despite the proposal's attractions (of which more shortly), the considerations just rehearsed seem decisively to refute it. For the facts revealed by attending to the scene before my eyes are at best very weak evidence that I see a hawk. Call this the *evidential objection* to the transparency proposal.

This chapter will elaborate and defend the transparency proposal. Let us begin by confirming that the evidential objection has no quick solution.

2.2. *Supplementary Environmental Evidence*

Transparency & the
evidential objection

According to the transparency proposal, I know that I see a hawk merely by attending to the hawk. And the evidential objection is that the evidence about the hawk gathered by this procedure is not good evidence that I see it. Notice, though, that vision also gives me information concerning the spatial relation between the hawk and myself, namely that I am facing the hawk.⁷ Now this does not help much on its own—I can easily face the hawk and not see it for the simple reason that I might have my eyes closed. But what if we add in evidence (provided by proprioception or kinesthesia) about the disposition of my eyelids, and other relevant bodily parts? Can't I then know that I see a hawk? Admittedly, vision is not then the only source of my evidence, but this revision preserves the basic idea that one knows what one sees by attending to one's environment, broadly construed.

Dretske in effect considers the revision and dismisses it in a few sentences: "Zombies, after all, have bodies too. . . . A zombie's arms and legs, just like ours, occupy positions. Their muscles get fatigued . . ." (2003, 2). In the skeptical scenario, the zombie's body also faces the hawk, the zombie's eyes are open, etc. This additional evidence does not discriminate, then, between the scenario in which I see a hawk and the scenario in which I am a (Dretskean) zombie.

Likewise, a video taken by a camera does not (usually) have the camera in it. Of course, one could work out the location and motion of the camera from a video containing enough perspectival information (cf. Gibson 1979, ch. 7), but the information the camera records about the scene does not itself imply that it was recorded. Although the emphasis will be on vision, the problem about the self will also be addressed in what follows, albeit inexplicitly.

⁶ For Dretske's earlier view, see Dretske 1995, ch. 2. The now numerous discussions of "transparency" in the literature are of course not sourced to Ryle [1949] 1963, but to Moore 1903.

⁷ That vision allows me to know something about myself is itself puzzling (see note 5). But let us grant it for the moment.

But even by Dretske's lights this is too quick. His question is: "What is it about this world [of "objective, mind-independent objects"] that tells you that, unlike a zombie, you experience it?" And his dismissal of the present proposal gives the impression that an answer needs to be absolutely skeptic-proof, displaying a body of evidence gained through perception that entails that I see a hawk. In fact, Dretske is not setting the bar so high: the challenge he poses is to explain how I know that I see a hawk by observing the environment (including, perhaps, my body). And the suggestion about proprioception and my relation to the hawk is, in effect, the idea that I find out that I see a hawk on the basis of the sort of evidence that would support the claim that someone else sees a hawk. I can come to know that someone else sees a hawk by noting that there is a suitably placed and salient hawk, that the person's eyes are open and converge on the hawk, and the like.⁸ Or so we may assume—skepticism about other minds is not the issue. Hence, the problem with the present suggestion is not that it fails to supply a way of knowing that I see a hawk, and so that I am not a zombie.

The problem, rather, is that I plainly do not rely on supplementary proprioceptive evidence in order to know that I see a hawk. Suppose that unbeknownst to me, I am suddenly stuck by a bizarre medical condition that renders my eyelids transparent. I turn to face the hawk and close my eyes. I know these latter two facts. Since I know that my eyes are closed, on the present suggestion I lack sufficient evidence to know that I see a hawk. But if Ryle asks me "Do you see a hawk?" I would hardly claim that I do not, or that the question was a difficult one to answer! I know that I see a hawk, just as I do in the normal case.⁹

3. Evans on Knowledge of Perceptual Experiences

The evidential objection would be of little interest if the transparency proposal lacked plausibility anyway. However, there is much to be said in its favor. As Ryle puts it, my seeing of the hawk "seems to be a queerly transparent sort of

⁸ Uncontroversially, I might sometimes have good evidence from vision that I see a hawk (perhaps I see myself in a mirror, staring at a hawk). But this is not a typical case.

⁹ Since I know that my eyes are closed, I do have some evidence that I do not see anything (and so must be hallucinating a hawk). But this contrary evidence surely does not prevent me from knowing that I see a hawk.

The point can be reinforced by considering other modalities, which one would expect to have the same basic epistemology as vision. Suppose I hear the distinctive scream of a red-tailed hawk and cannot identify the direction of the sound. I can know that I hear the scream (or the hawk) without checking that my ears are not blocked, or gathering further evidence about the location of the hawk and the orientation of my ears.

process.”¹⁰ There is the hawk, sitting on the fence post. There is Gilbert Ryle, out for a stroll, pausing to descry the hawk. Here am I, looking at Ryle and the hawk. To me, Ryle’s seeing of the hawk is a perceptually manifest fact, as is the fact that the hawk is on the fence post. My own seeing of the hawk, on the other hand, is quite a different matter. I see Ryle and note that his gaze is hawkwards; I do not see myself, or my eyes. Moreover, **it does not ring true to say that I discover that I see the hawk by some special introspective sense. There is no switch in attention—say to myself or to a “visual experience”—when Ryle asks me “Do you see the hawk?” I answer by attending to the hawk. (Indeed, if I attend to something else, I might well give the wrong answer.**¹¹)

Evans on transparency

The transparency proposal can be extracted from Evans’s influential but brief discussion of the “self-ascription of perceptual experiences” (1982, 226):

[A] subject can gain knowledge of his internal informational states [his “perceptual experiences”] in a very simple way: by re-using precisely those skills of conceptualization that he uses to make judgements about the world. Here is how he can do it. He goes through exactly the same procedure as he would go through if he were trying to make a judgement about how it is at this place now . . . he may prefix this result with the operator ‘It seems to me as though . . .’ (ibid., 227–228)

Here Evans is concerned with knowledge of how things perceptually appear. But the point is evidently supposed to apply to knowledge of what one sees. The subject—Evans might have added—may, after looking at the scene before his eyes, prefix a phrase encapsulating the result (‘a hawk’) with ‘I see’.

The transparency procedure

Although the quotation has the subject attaching a sentential operator to a sentence, presumably Evans did not mean to tie knowledge of one’s perceptual states to language. **Recast in nonlinguistic terms and restricted to the case of seeing an object, the procedure suggested by the quotation is that one can come to know that one sees an object by an inference whose sole premise concerns one’s (typically non-mental) environment, “how it is at this place now,” as Evans puts it.** (In fact, this is not Evans’s view. This will become clear later, in section 8, when the elision in the quoted passage from Evans is filled in.)

If we remain similarly coy for the moment about the exact nature of the premise, this inference can be set out as follows:

¹⁰ Unfortunately, Ryle then goes on to claim that “the mystery dissolves when we realize that ‘see’, ‘descry’, and ‘find’ are not process words, experience words, or activity words. . . . The reason why I cannot catch myself seeing. . . is that [this verb is] of the wrong type to complete the phrase ‘catch myself. . .’ ([1949] 1963, 285). Since the mystery can be stated without falsely assuming that ‘see’ is a “task verb” (Ryle’s phrase) like ‘run’ and ‘aim’, Ryle’s proposed solvent does not work.

¹¹ With the defensible assumption that one may see an object without attending to it.

p: It is thus-and-so at this place now

c: I see a hawk

With this more explicit statement of the transparency proposal in hand, Dretske's evidential objection can be put as follows: this inference is not knowledge-conducive because p is not good evidence for c.

3.1. Epistemic Rules

For the discussion to come, some terminology will be useful. Let us say that an *epistemic rule* is a conditional of the form, 'If conditions C obtain, believe that p', for example:

Epistemic rules

WOODPECKER: If x is a bird with a red head, believe that x is a woodpecker.

And let us stipulate, not unnaturally, that one *follows* this rule on a particular occasion iff one believes that x is a woodpecker because one recognizes that x is a red-headed bird, where the 'because' marks the kind of reason-giving causal connection that is discussed under the rubric of 'the basing relation'.¹²

basing

In general, then, S follows the rule 'If conditions C obtain, believe that p' on a particular occasion iff on that occasion:

1. S believes that p because she recognizes that conditions C obtain; which implies:
2. S recognizes (hence knows) that conditions C obtain;
3. conditions C obtain;
4. S believes that p.

Following WOODPECKER (in certain circumstances that can be left uninvestigated) tends to produce knowledge, and hence is a *good* rule. Following DODO, 'If x is a quacking bird, believe that x is a dodo', produces beliefs that are not knowledge, and hence is a *bad* rule.

The Evans-inspired transparency proposal and Dretske's evidential objection can be put in terms of this apparatus of epistemic rules. On the transparency proposal, I come to know that I see a hawk by following this rule:

HAWK: If it is thus-and-so at this place now, believe that you see a hawk.¹³

And Dretske's evidential objection is that if 'it is thus-and-so' is spelled out as intended, as concerning the hawk before my eyes, the rule must be a bad one, because the antecedent is very poor evidence that I see a hawk.

¹² In the terminology of Byrne 2005, 94, because of the presence of the schematic letter 'x', WOODPECKER is a *schematic* rule; one follows a schematic rule iff one follows an instance of it.

¹³ 'You' refers to the rule-follower; 'now' and 'this place' refer to the time and place the rule is followed.

In fact, Dretske's objection can arguably be overcome, but only if another objection is dispatched first. Explaining this is the burden of the next two sections.

4. Knowing What One Knows and the Evidential Objection

Often one knows what one knows. I know that I know that there is a hawk on the fence post, for example. How do I know that? Adapting a famous passage from Evans on the epistemology of belief: if someone asks me, "Do you know that there is a hawk on the fence post?" I must attend, in answering him, to precisely the same outward phenomena, as I would attend to if I were answering the question "Is there a hawk on the fence post?"¹⁴ This suggests that I know that I know that there is a hawk on the fence post by following the rule:

KNOW: If p , believe that you know that p .

Whether or not this suggestion is correct, it should not be in dispute that we at least have the capacity to follow **KNOW**: often we know that p , and are capable of performing elementary inferences.

Now Dretske's evidential objection arises for **KNOW**. The fact that there is a hawk on the fence post is hardly good evidence that I know that there is a hawk on the fence post. Hence, if I follow **KNOW** and infer that I know that there is hawk on the fence post from the premise that there is hawk on the fence post, then my conclusion is based on inadequate evidence, and so is not known.

KNOW is self-verifying

Notice, though, that **KNOW** has the interesting property that, necessarily, if one follows it, one's resulting belief about one's state of knowledge is true. This is because one can only follow **KNOW** if one *knows* that p . **KNOW**, as we can put it, is *self-verifying*. In this respect, **KNOW** is unlike typical good rules that we follow in ordinary life, such as **WOODPECKER**. **WOODPECKER** is not self-verifying because one may know, of a certain non-woodpecker (a red-crested cardinal, for instance), that it is a red-headed bird.

That **KNOW** is self-verifying is enough to blunt the force of the evidential objection. Usually, if one reasons from inadequate evidence, then one's conclusions will

¹⁴ For the passage (about "a third world war"), see Evans 1982, 225. The passage overreaches slightly: Evans says that I can answer "the question *whether* I believe that p " (my emphasis) by this procedure, which is incorrect in cases where I have no opinion either way. The point also applies to knowledge: if I am clueless about the location of nearby hawks, Evans's procedure does not apply. Yet I can easily know that I do not know that there is a hawk on the fence post.

be false. And if they are true, that will be by accident, and hence they will not amount to knowledge. But the beliefs produced by following KNOW are true, and nonaccidentally so.¹⁵ Admittedly, this is not sufficient for these beliefs to amount to knowledge, but why think they do not? If in these special cases reasoning from inadequate evidence results in beliefs with one of the characteristic signatures of knowledge, then the status of the evidential objection is moot.

Return to Dretske's evidential objection to the goodness of:

HAWK: If it is thus-and-so at this place now, believe that you see a hawk.

The upshot of the comparison with KNOW is that if HAWK, like KNOW, yields nonaccidentally true beliefs, then the evidential objection is muted, if not entirely silenced. But does following HAWK lead to nonaccidentally true beliefs? To answer that question the template 'thus-and-so at this place now' needs to be filled in. And as soon as we try to do that, another—potentially more serious—objection is apparent.

5. The Amodal Problem

A first thought is to fill in the 'thus-and-so' along these lines:

HAWK†: If there is a hawk over there, believe that you see a hawk.

On second thought, HAWK† is not a good rule, even waiving Dretske's evidential objection. Suppose I follow HAWK†, and so know that there is a hawk there. Since there are numerous ways of knowing that there is a hawk there that do not involve currently perceiving the hawk, let alone seeing it, the probability that I see a hawk, given that I know that there is a hawk there, is low. To conclude that I see a hawk is to take a stab in the dark.

Suppose we try inserting the subject into the antecedent:

HAWK†: If there is a hawk right in front of you, believe that you see a hawk.

This sort of maneuver certainly helps increase the probability that I see a hawk, conditional on my knowing the antecedent. But again, there are many other non-visual ways in which I might know the antecedent. If following a rule like HAWK† were my chief strategy for finding out that I see a hawk, then I would be prone to all sorts of errors that I actually never make.

The apparent root of the difficulty is that information does not wear its provenance from a particular sensory modality on its face—information is amodal. Perhaps

¹⁵ Cf. Byrne 2005, 96–98.

the idea that one knows what one believes and knows by directing attention “outward—on the world” (Evans 1982, 225) has something going for it. But the amodal nature of information, it might be thought, shows that *perception* is where this idea irretrievably runs into sand.

What are the alternatives?

6. Alternatives to Transparency

According to the transparency proposal, I know that I see a hawk by an inference from a single premise about the hawk-infested landscape beyond. There two main alternative options.

Option 1 is that no premise about my environment is needed: I know that I see a hawk without appealing to evidence concerning the scene before my eyes—I know *non-observationally* that I see a hawk. (‘Observation’ is meant to cover only observation of the normal visual sort; on option 1, perhaps I know that I see a hawk by introspection, conceived of as a kind of inner perception.) Option 2 is that although a premise about my environment is needed, it is not enough: additional mental evidence is required. Let us take these in turn.

6.1. Option 1: Non-Observational Knowledge

Option 1, that I know non-observationally that I see a hawk, requires immediate amendment. First, note that this does not apply to every case of knowing that I see a hawk, because sometimes an environmental premise is plainly needed: I know that I see that bird (pointing to a hawk perching atop a distant tall tree), but I am not in a position to know that I see a hawk. Ryle is passing by and informs me that the bird is a hawk; with this environmental premise in hand, I conclude that I see a hawk. Second, extending this first point, perhaps one can *never* know non-observationally that one sees *a hawk*—all such knowledge is based on evidence that one sees such-and-such, and that such-and-such is a hawk, with the latter item of evidence being known observationally. So a more careful and general statement of option 1 is as follows: knowledge that one sees an F/*this F* is either non-observational, or else based on evidence that includes the fact that one sees a G/*this G*, known non-observationally.

If there is any non-observational knowledge of this sort, knowledge that one sees this red spot (pointing to a clearly visible red spot) is an example, or so we may suppose. Since the fact that one sees this red spot entails that this spot is red, one may come to know that this spot is red by inference from the fact that one sees this

red spot. Now one may also know that this spot is red simply by looking at it—an animal with no conception of seeing could use its eyes to know that this spot is red. So no knowledge that one sees this red spot is necessary. Thus, on this view, there are two routes to the same conclusion: one may know that the spot is red twice over, by inference from a non-observationally known fact about what one sees, and by the more familiar method of simply using one's eyes.

This result is more than strange. First, note that one may see what is, in fact, a red spot, even though the spot does not look red (perhaps one is viewing the spot in very dim light). One is not able to tell by looking that this spot is red, but one might have various backup routes to that conclusion—perhaps one painted the spot oneself from a can of red paint. However, the alleged non-observational backup route is clearly inoperative: although it is true that one sees this red spot, no amount of introspection will reveal this fact. The obvious explanation is that in cases where the spot is clearly visible, the information one obtains by vision about the spot is somehow used to derive the conclusion that one sees this red spot, but if that is right then option 1 must be rejected.

Second, note that when one sees a red spot and believes both that this spot is red and that one sees this red spot, it is not a possibility that two spots are in play. Could this red spot be a different spot from this red spot that one sees? That is not a serious question, but if one knew non-observationally that one sees this red spot it apparently would be. Return to the situation in which one views this red spot in dim light. Suppose one remembers that one painted this spot red; on occasion, one might reasonably wonder whether one's memory was quite accurate—perhaps one painted another spot red, not this very spot. As before: the obvious explanation of why the identity of the spot is never in question is to say that the information about this spot is used to derive the conclusion that one sees it.

Finally, if I know non-observationally that I see this red spot, then certain dissociations are to be expected. In particular, one's vision and reasoning capacities might be working perfectly normally, while the mechanism that yields non-observational knowledge that one sees this red spot is broken or absent. One's only means of finding out that one sees a red spot would then be similar to third-person cases: one knows that one sees a red spot because one knows that there is a red spot right there, that the light is good, that one's eyes are open, and so forth. Often one knows through vision about an object's location and other features, but is unsure whether someone else sees it (perhaps one does not know that the person's gaze is in the right direction). Similarly, someone who only had third-person access to her states of seeing would sometimes be in a state of uncertainty about whether she saw an object, while quite certain (via her excellent vision) about the nature of the object itself. It is safe to say that this bizarre condition

never occurs.¹⁶ Pending some explanation of why the non-observational mechanism never fails in this way, this is a reason for thinking that option 1 is incorrect.¹⁷

6.2. Option 2, First Pass: Visual Sensations

Since option 1 faces some serious objections, let us turn to option 2, that additional evidence is required. And from a more traditional position in the philosophy of perception, the need for such evidence is palpable. Seeing an object is a matter of the object causing distinctive sorts of affectations of the mind, “visual sensations.” It is thus natural to think of knowledge that one sees an object as resting on evidence about both ends of this causal transaction—evidence about the object coming from observation, and evidence about the sensation coming from some other source. So,

¹⁶ As Ned Block pointed out to me, the closest approximation in the literature appears to be the case of “reverse Anton’s syndrome” described by Hartmann et al. (1991; see also Block 1997, 159). The patient was initially diagnosed as blind due to a stroke. Two years later he was found to have spared vision in a 30° wedge in both fields. Anton’s syndrome patients deny that they are blind; this patient denied that he could see. At one point he remarked that “you (the examiners) told me that I can see it, so I must be able to see it” (Hartmann et al. 1991, 33). However, the patient’s vision was far from excellent. He could read words, but with limited accuracy (51% correct on a standard test). Strikingly, he was “unable to discriminate light from dark” (37). The patient’s cognition was also impaired, with mild language and memory deficits. Further, sometimes he used perceptual verbs in describing his condition: on a color-naming task, “he maintained that he could “feel” or “hear” the color” (34). The correct description of the patient’s predicament is unobvious. As Hilbert notes, “a certain amount of scepticism about the case is in order” (1994, 449). And, as Hartmann et al. say, reverse Anton’s syndrome is not clearly documented in any other published case.

¹⁷ The bizarre condition is what Shoemaker (1994) calls “self-blindness” with respect to seeing. It is worth emphasizing that someone who is self-blind with respect to vision is not the “super-duper blindsighter” of Block, who has “blindsight that is every bit as good, functionally speaking, as [normal] sight” (1997, 409). There are two differences. The first is that Block’s super-duper-blindsighter denies that he sees anything, whereas the self-blind person knows by third-person means that she sees things. This first difference probably just reflects unimportant differences of detail between these two science-fiction stories; Block would not deny that a super-duper-blindsighter could investigate his own states of seeing third-personally.

The second difference is the important one: the super-duper blindsighter has a faculty “that is every bit as good, functionally speaking,” as normal vision, except that the resulting perceptual states lack “phenomenal consciousness.” The super-duper-blindsighter is thus, as Block says, a “quasi-zombie” (1997, 409), or a “visual-zombie,” in something close to the usual sense of ‘zombie’ (not Dretske’s sense: see section 2). The self-blind person has normal vision, at least in the sense that she sees what we see, and has the sort of perceptual knowledge that we have, but lacks the “peculiar access” (Byrne 2005) that we have to our states of seeing. Nothing is being assumed, one way or the other, about whether the self-blind person’s visual states have phenomenal consciousness in Block’s sense.

Dissociation problems also afflict option 2, but this will not be discussed further.

Self-blindness vs.
blindsight

in order to know that I see a hawk, I need to know, *inter alia*, that I am having a visual sensation. Such a sensation is an occurrence in my mind, not on the fence post beyond, so no wonder peering at the hawk is not sufficient.

An analogy can clarify the traditional position further. I am holding a nettle, and feel a stinging pain in my hand. How do I know the additional fact that the nettle is stinging me (i.e., causing the pain)? It would be a mistake to investigate the issue by concentrating solely on the nettle; rather, I need to attend to something else entirely, namely the pain in my hand. Putting these two items of evidence together—that I am holding a nettle and that I have a pain in my hand—I can conclude that nettle is stinging my hand. That conclusion is not entailed by my evidence, but in the circumstances my evidence strongly supports it. Likewise, **on the present suggestion, I can conclude that I see a hawk on the basis of two items of evidence: the external nonpsychological fact that a hawk is present and the internal psychological fact that a visual sensation is occurring.** (Note that placing a substantive restriction on the type of visual sensation would not be advisable, since almost any kind of visual sensation could accompany seeing a hawk—it could look blue, or cubical, or whatever.) Knowledge that one sees an F, then, is obtained by following this rule:

SEEⁱ: If an F is present and you are having a visual sensation, believe that you see an F

Evidently SEEⁱ is hopeless. Taking the existence and epistemology of “visual sensations” for granted, on many occasions one knows that an F is present and that one is having a visual sensation, yet one does not see an F. SEEⁱ is thus a bad rule. Moreover, we do not follow it. Suppose I see a sheep in a field; although no hawk is in sight, I know that there is a hawk in the vicinity. I have no inclination to follow SEEⁱ and conclude that I see a hawk. Although rules are generally defeasible—despite knowing that *x* is a red-headed bird, one might have additional evidence that prevents one following WOODPECKER—it is unclear what the defeater might be in this case.¹⁸

We can pass over attempts to add epicycles to SEEⁱ, because the nettle analogy is fundamentally defective. When I see a hawk I do not have a spectacular kind of migraine headache whose only connection to the hawk is that it is caused by the hawk. This is basically Ryle’s point when he observes that in the “unsophisticated use of ‘sensation’” a typical case of seeing does not involve any sensations ([1949]

¹⁸ Another problem is due to the word ‘present’ in the antecedent. This prevents me from always believing that I see a hawk, since I always believe that there are hawks somewhere. But ‘present’ excludes too much—in principle, I can see a hawk at any distance (cf. seeing a supernova) and also readily know that I see it.

1963, 228). One can know what stinging sensations are without knowing anything about nettles, but insofar as the philosophical notion of a “visual sensation” is intelligible, it is not likewise only externally related to its causes. **Visual sensations or, better, visual experiences, are specified in terms of the region of the external world that they purportedly reveal.** That is, when I look at the hawk and recognize it as such my visual experience is an experience *of a hawk*. Does this reconception of visual sensations as visual experiences help rescue Option 2?

6.3. *Option 2, Second Pass: Visual Experiences of an F*

Start by applying the reconception to SEEⁱ:

SEEⁱⁱ: If an F is present and you are having a visual experience of an F, believe that you see an F.

This straightforwardly copes with the case where I see a sheep in a field and know that there is a hawk in the vicinity, which I do not see. I do not have an experience of a hawk, and so am not in a position to follow SEEⁱⁱ.

But what is it for a visual experience to be “of” a hawk? An influential discussion of this question is in Searle’s 1983 book *Intentionality*. Searle writes:

I can no more separate this visual experience from the fact that it is an experience *of a yellow station wagon* that I can separate this belief from the fact that it is a belief that it is raining; the “of” of “experience of” is in short the “of” of Intentionality. (Searle 1983, 39)

An experience of a hawk may be said to be “of” a hawk in the same way that a belief about a hawk is “of” a hawk. Experience, then, like belief, has intentionality: my experience of a hawk and the belief that there is a hawk on the fence post are both “of” or “about” a hawk. But the parallel, Searle thinks, is even closer. The belief that there is a hawk on the fence post has propositional content, namely the proposition that there is a hawk on the fence post. And likewise for visual experiences:

The content of the visual experience, like the content of the belief, is always equivalent to a whole proposition. Visual experience is never simply *of an object* but rather it must always be *that such and such is the case*. (40)

In the case of an experience of a yellow station wagon, “a first step in making the content explicit,” Searle says, “would be, for example, I have a visual experience (that there is a yellow station wagon there)” (41).¹⁹

¹⁹ Searle’s considered view is that the content is the proposition that “there is a yellow station wagon there and that there is a yellow station wagon there is causing this visual experience” (1983, 48), which has attracted a lot of criticism. See, e.g., Recanati 2007, ch. 17.

The content of
experience

On this view, perceptual experiences have content, like belief, desire, and other “propositional attitudes.” To a first approximation, one may think of the content of the subject’s visual experience as the information (or misinformation) delivered to the subject by his faculty of vision (cf. Armstrong 1968: 224). When the delivery is one of misinformation, the subject suffers a visual illusion. Although this is somewhat controversial, it is at least a huge advance over the sense datum theory, and the traditional view mentioned in the previous section.²⁰

Assume, then, that visual experiences have contents, *v-propositions*; true *v-propositions* are *v-facts*.²¹ Let ‘[. . . F(x) . . .]_v’ be a sentence that expresses a particular *v-proposition* that is true at a world *w* only if *x* is *F* in *w*. Read ‘You V[. . . F(x) . . .]_v’ as ‘You have a visual experience whose content is the proposition that [. . . F(x) . . .]_v’. Then a more explicit version of SEEⁱⁱ is:

SEEⁱⁱⁱ: If you V[. . . F(x) . . .]_v and *x* is an *F*, believe that you see an *F*.

What exactly are *v-propositions*? Searle’s example—the proposition that there is a yellow station wagon there—is at best a “first step,” as he says: it hardly begins to capture the apparent scene before the eyes when one sees a yellow station wagon. In fact, it might not even be a first step. Does the content of visual experience concern station wagons, hawks, and the like, as such? If the ostensible yellow station wagon is actually white, vision is surely to blame for delivering misinformation to the subject. But what if the ostensible station wagon is a sedan? Here there is a temptation to exonerate vision, and instead to point the finger at the subject’s judgment that the car is a station wagon. The issue is less than clear, and in any event disputed.

Granted that visual experiences have contents, it is not disputed that the content at least concerns what falls under the rubric of “mid-level vision” in vision science: shape, orientation, depth, color, shading, texture, movement, and so forth: call these *sensible qualities*. In fact, without begging any important questions we can restrict *v-propositions* so that they just concern sensible qualities. (This restriction is imprecise, but that will not matter here.) With this restriction, and letting ‘[. . . x . . .]_v’ express a *v-proposition* that is true at a world *w* only if *x* has a certain sensible quality in *w* (i.e., if *x* is red, or square, . . .), we get:

SEE^{iv}: If you V[. . . x . . .]_v and *x* is an *F*, believe that you see an *F*.

²⁰ For skepticism about the existence of “visual experiences” as Searle and many other philosophers conceive of them, see Byrne 2009. That skepticism will be suspended here.

²¹ Some disjunctivists will deny that illusory cases involve having a visual experience whose content is a false proposition (see the introduction to Byrne and Logue 2009). However, they have no special reason to dispute this chapter’s account of the epistemology of veridical cases.

V-facts'

Notice that because 'F' does not appear in the scope of 'V', this is an improvement on SEEⁱⁱⁱ. Return to an example given in section 6.1: I see a bird atop a tall tree, too far away to make out its shape or color, which Ryle tells me is a hawk. I am presumably not having an experience "of a hawk," since the information available to my visual system is too impoverished. I therefore cannot follow SEEⁱⁱⁱ. But I can follow SEE^{iv}, since I am having a visual experience with a content that concerns the hawk, albeit a content that does not identify it as such.

Although SEE^v is the best attempt so far, it is not good enough. Recall that '[. . . x . . .]_v' expresses an object-dependent proposition—one whose truth at world *w* depends on how a certain object (namely *x*) is in *w*. Further, it is very plausible that one can only enjoy a visual experience with such an object-dependent content in a world in which the object exists (at some time or another). That is, 'YouV[. . . x . . .]_v' entails 'x exists'. But if that is right, then we are back in the same bind that afflicted the suggestion that I can know that I see a hawk non-observationally (see section 6.1). Since the existence of *x* is entailed by the proposition that IV[. . . x . . .]_v I have two routes to the conclusion that that object (the hawk) exists. And, as before, the account leaves open a possibility that should be closed, namely that there are two objects, one known about through vision, and the other known about through non-observational means.

Can these problems be avoided by denying that *v*-propositions are object-dependent? The view is not well motivated. By perceiving, in particular by seeing, one may come to know things about individual objects in one's environment—that *that* is a hawk, for example. It is thus natural to think that the information delivered by vision is object-dependent: the testimony of one's visual system concerns this very hawk. Still, this alternative needs examining further.

Suppose, then, that when I see the hawk, it is not pinned down by a *v*-proposition with the hawk as a constituent, but rather by a proposition that identifies the hawk by description. (For the sake of the argument, we can ignore the difficult question of what this description exactly is.) Here is the descriptive counterpart of SEE^{iv}:

SEE^v: If you V[. . . (the G) . . .]_v and the G is an F, believe that you see an F

Apart from paucity of motivation, is there anything wrong with it?

Consider a case where I think or suspect that I am suffering from an illusion. I see a hawk, but I doubt that the hawk is the way it looks. Perhaps the hawk looks like a penguin right in front of me, and I have reason to believe that this is the product of a devious arrangement of distorting mirrors, with the ordinary-looking hawk being positioned behind my back. However 'the G' is filled in, we may safely suppose that I do not know or believe that the hawk is the G. SEE^v is thus of no help. Nonetheless, nothing prevents me from knowing in an ordinary sort of way that I

see a hawk. For instance, I know that I see a hawk because Ryle tells me that *this* (clearly referring to the penguin look-alike before me) is a hawk. Since there is nothing epistemologically special about this case, if SEE^v does not explain my knowledge here, it does not explain it elsewhere.

Even taking the ontology and epistemology of "visual experiences" for granted, there are no easy alternatives to the transparency proposal. So let us revisit it.

7. Back to Transparency

For the moment, shelve illusions and concentrate on veridical cases, where one sees an object and it is as it looks. Return to the object-dependent suggestion:

SEE^v: If you V[. . . x . . .]_v and x is an F, believe that you see an F

The problems just rehearsed are all in effect traceable to the 'V', which suggests the experiment of dropping it. And removing the 'V' yields a version of the transparency proposal:

SEE: If [. . . x . . .]_v and x is an F, believe that you see an F

The amodal problem of section 5 seemed to doom the transparency proposal. Could v-facts rescue it?

Recall that a v-fact concerns the sensible qualities of objects in the scene before the eyes. In one way this notion is perfectly familiar. When I see the hawk on the fence post, a segment of the visible world is revealed: an array of colored, textured, three-dimensional objects, casting shadows, some occluding others, at varying distances from my body, with various illumination gradients, and so forth. A certain v-fact just specifies that array, the scene before my eyes. If Ryle and I strike up a conversation about the spectacular view of the North York Moors, v-facts are our topic.

On the other hand, giving a theoretically satisfying characterization of v-facts is difficult. Armstrong, for instance, speaks of perceptual content as comprising "certain very complex and idiosyncratic patterns of information about the current state of the world" (1968, 212), while declining to be much more specific.²² Even vision science often in effect dodges the issue with placeholders like 'visual representation'. Complexity or informational richness is no doubt part of the story, but even in the case of viewing a very simple scene—say, a red spot against a grey background—it is unclear how to proceed. Just concentrating on one feature of the

²² A rare example of a more detailed account is in Peacocke 1992.

spot, its hue, the predicate 'is red' (or even some made-up predicate like 'is red₂₉') does not quite do it justice. The particular red hue of the spot might be a little yellowish, or alternatively a little bluish; how exactly information about the hue is packaged by vision is not at all obvious.²³ Even though the familiar may resist theory, fortunately for our purposes not much theory is required.

Vision, we may say, reveals the *visual world*: the world of *v-facts*. In the visual world things are colored, illuminated, moving, and so on, but not smelly or noisy.²⁴ Likewise, olfaction reveals the *olfactory world*: the world of *o-facts*. The olfactory world—at least, our olfactory world—is a relatively impoverished place, consisting of odors located around the perceiver's body. The auditory world, the world of *a-facts*, is considerably more complicated, consisting, inter alia, of sounds of varying loudness and pitch at a variety of locations.

One may base one's actions and inferences on how things are in the visual world—this just requires a sensitivity to different aspects of one's environment. (In particular, it does not presuppose self-knowledge.) Suppose one investigates one's environment and finds that a certain *v*-fact, the fact that $[\dots x \dots]_v$ obtains. Vision is, at least in creatures like ourselves, an exclusive conduit for *v*-facts. Hence, one's information source must be vision, not audition, olfaction, testimony, or anything else. Although information is amodal in principle, for us *v*-facts do indicate their provenance—(visual) information is *practically* modal. Thus, SEE apparently solves the amodal problem.

What about the evidential problem? That has not gone away, because the fact that $[\dots x \dots]_v$ remains stubbornly devoid of vision. That is, the hawk before my eyes, with its rich variety of visual sensible qualities, offers no indication at all that it is seen.

Still, SEE takes the sting out of the evidential objection much as KNOW did. Recall that the latter rule is:

KNOW: If *p*, believe that you know that *p*.

Section 4 noted that KNOW is self-verifying: if one follows it, then one's belief that one knows that *p* is true. KNOW typically produces reliably true beliefs, and there is no clear barrier to supposing that it also typically produces knowledge. To that extent, the evidential objection is rebutted.

²³ For a sketchy proposal about the visual representation of hue, see Byrne and Hilbert 2003, 14.

²⁴ The last two sentences should be qualified. First, cross-modal effects show vision does not reveal the visual world unaided—other modalities sometimes help too. Second, the negation takes wide-scope: it is not the case that there are smells or noises in the visual world. The visual world is silent on such matters.

Practical self-verification & the evidential objection

SEE, in contrast, is not self-verifying: perhaps one could in principle learn that $[\dots x \dots]_V$ by reading it in the—as-yet-unwritten—language of vision; one would not thereby see x . But it is *practically* self-verifying: in all ordinary situations, one knows that $[\dots x \dots]_V$ only if one sees x . As section 4 also noted, self-verification is not sufficient for a rule to be knowledge-conducive; a fortiori, practical self-verification is not either. But the dialectical imperative is not to prove that following SEE yields knowledge; rather it is to reply to the evidential objection. And for that, practical self-verification will do.

7.1. The Memory Objection

The claim that SEE is practically self-verifying might be thought to be too strong. Surely, if v -facts can be known, they can be remembered. Shouldn't we then have said: in all ordinary situations, one knows that $[\dots x \dots]_V$ only if one sees or saw x ? And if so, there is the following difficulty.

Suppose I see a red spot (x) at time t_1 . Write the relevant v -fact as 'the fact that $[\dots \text{Red}(x, t_1) \dots]_V$ ', and further suppose that I remember it. Shortly after, at t_2 , a piece of cardboard is placed in front of the spot, completely occluding it; I am quite confident that the spot itself has not changed color: the distinctive visual way the spot was is the way the spot now is. I know (we may assume) that $[\dots \text{Red}(x, t_2) \dots]_V$. Granted all this, I am in a position to follow SEE and conclude that I see a red spot. But obviously I do not. Why not? Either something blocks in the inference in this case, or I do not follow SEE in any circumstances.

Once this disjunction is conceded, it is hard to avoid the second disjunct.²⁵ However, the memory objection should not be allowed to get started in the first place. Consider recalling something one has seen—say, recalling the red spot one saw. Such an exercise of one's episodic memory is akin to visual imagery: it is somehow visual in character, but easily distinguished from actually seeing a red spot. In a Humean framework, this is because, in episodic recollection, one is aware of a faded copy (an "idea") of one's past visual "impression" (or sense datum). The similarity is explained by the copying, the difference is explained by the fading. Granting the existence of impressions and ideas, this purported explanation is perhaps the best that can be found.

²⁵ Could the fact that the cardboard "occludes" the spot block the inference? No. If 'occludes the spot' means 'prevents me from seeing the spot', this just raises the question how I know the cardboard occludes the spot. On the other hand, if it means 'is opaque and in front of the spot' then my knowing this fact does not explain why I do not follow SEE. Suppose I can in fact see the spot, due to some devious arrangement of mirrors, or because I have suddenly gained Superman's ability to see through walls. Despite knowing that the cardboard is opaque and in front of the spot, I would follow SEE and conclude that I see it.

The Humean attempt at an explanation has a (superior) counterpart in the present information-based framework. In episodically recollecting the spot, one is aware of a segment of the past visual world, but although the information is packaged visually, it is a transformed and degraded version of the visual information that characterizes successful seeing. The similarity is explained by the visual packaging, the difference is explained by the transforming and degrading.²⁶

If that is right, then it was too hasty to say that I remember that $[\dots \text{Red}(x, t) \dots]_v$. What I do remember could be written as 'the fact that $\{\{\dots \text{Red}(x, t) \dots\}_v\}$ ', where the curly brackets represent the episodic-memory transformation and degradation of the true v -proposition that $[\dots \text{Red}(x, t) \dots]_v$. And if we not unreasonably assume that memorially transformed and degraded v -facts are disjoint from v -facts, then remembering that $\{\{\dots \text{Red}(x, t) \dots\}_v\}$ will not thereby put me in a position to follow SEE, and the memory objection fails.²⁷

On this account, one might expect that in some cases, where other cues are absent, episodic visual recollection (or visualizing more generally) will indeed be mistaken for seeing. Since this is basically the converse of the Perky effect (where one mistakes seeing for visualizing), it can be called the *converse-Perky effect*. And there is some evidence that the converse-Perky effect occurs.²⁸

8. Evans Again, and the Known-Illusion Problem

So far we have concentrated on the veridical case: I see the hawk and it is as it looks. Let us now return to illusions. To give some examples more realistic than the one mentioned at the end of section 6: the hawk looks closer than it really is, or a shadow appears as a patch of darkened green on the field beyond, or the hawk is perching on a wall that generates Richard Gregory's "café wall illusion."²⁹ In such cases, the fact I seem to apprise, that $[\dots x \dots]_v$ where x = the hawk, is no fact

²⁶ For a discussion of this, see Byrne 2011.

²⁷ Our visual memories are very impressive, at least under some conditions (Brady et al. 2008). But as far as I know, the evidence supports the transformation and degradation hypothesis. Some examples: Burnham and Clark 1955 (memory for hue), Uchikawa and Ikeda 1986 (memory for brightness). Clearly much more could be said here, though.

²⁸ See Goldenberg, Müllbacher, and Nowak 1995 (blindness denial arguably explained by the patient's spared visual imagery; see also Byrne 2010, 117–118). A more common example might be this. Suppose one is in bed, in almost total darkness. One opens one's eyes and looks in the direction of a familiar object—a desk, say. One episodically recollects its distinctive visual qualities as they appear from this angle. Does one see the desk? When in this sort of situation myself, I sometimes wonder whether I am merely visually recollecting the object rather than seeing it.

²⁹ See <http://en.wikipedia.org/wiki/Café_wall_illusion>.

at all. Still, I can easily discover that I see a hawk, just as I did in the original veridical example.

If I do not know that I am illuded, this case presents no difficulty. Say that one *tries to follow* the rule ‘If conditions C obtain, believe that *p*’ iff one believes that *p* because one *believes* that conditions C obtain. If one follows a rule one tries to follow it, but not conversely.³⁰ One cannot follow SEE if the relevant *v*-proposition is false, but one can try to follow it. And in the illusory example of the previous paragraph, if I try to follow SEE, then I will end up with a nonaccidentally true belief that I see a hawk, for essentially the same reasons as before.³¹

The problem, rather, is similar to the one faced by SEE’ at the end of section 6, and concerns the case when I know (or believe) that I am illuded. The method I use to discover what I see does not obviously alter when I know (or believe) that the hawk is not the way it looks: I can still know that I see it by attending to the hawk. If the transparency proposal applies at all, it must apply unmodified across the board. But if I do not believe the relevant *v*-propositions, I cannot even try to follow SEE. Hence, cases of known-illusion threaten to blow the transparency proposal entirely out of the water.

8.1. Evans’s Proposal

Recall the quotation from section 3, where Evans is explaining how someone may gain knowledge of how things perceptually appear by “re-using precisely those skills of conceptualization that he uses to make judgements about the world.” The quoted passage contained an elision, and it is time to restore it. Here are the crucial sentences.

[The subject] goes through exactly the same procedure as he would go through if he were trying to make a judgement about how it is at this place now, but excluding any knowledge he has of *an extraneous kind*. (That is, he seeks to determine what he would judge if he did not have such extraneous information.) (1982, 227–228)

Consider the following case. I am staring at what I know to be a gray patch on a green background. Due to a color contrast effect, the patch will look slightly

³⁰ See Byrne 2005, 97.

³¹ An objection at this point is that one cannot come to know something by inference from a false premise, a moral commonly drawn from Gettier cases. But an arguably better diagnosis of the Gettier cases is that *safety* (in the sense of Sosa 1999 and Williamson 2000, ch. 5) is a necessary condition for knowledge, not that no reasoning through false steps is a necessary condition for knowledge. And beliefs produced by trying to follow SEE will often be safe (cf. Byrne 2005, 96–98). See also Silins, this volume.

reddish. Since I am aware of the effect, I do not believe the relevant *v*-proposition, that $[\dots \text{Reddish}(x) \dots]_v$, where x = the patch. I know that I see a grey patch, but cannot know this by trying to follow SEE.

Evans's remarks suggest the following two-step alternative. First, I verify a certain counterfactual truth: if I had not known extraneous facts, I would have judged that $[\dots \text{Reddish}(x) \dots]_v$. That tells me that x looks reddish, and so that I see x . I then add in the fact that x is a grey patch, and conclude from this that I see a grey patch.

One immediate problem with this suggestion turns on the notion of "knowledge of an extraneous kind."³² The effect of excluding extraneous knowledge is intended to make me rely exclusively on the testimony of vision, but it cannot be characterized as "facts I know other than by current vision" on pain of circularity. Could an extraneous piece of knowledge be characterized simply as something that I previously knew about the patch? Then the counterfactual to be verified is 'If I had not known anything about the patch beforehand, I would have judged that $[\dots \text{Reddish}(x) \dots]_v$ '. This suggestion has a number of problems. First, it is quite implausible that a counterfactual of this sort will always be true in every case, or that I will judge that such a counterfactual is true.³³ Second, intuitively it gets things back to front. If I do know that the counterfactual is true, then isn't this because I know the patch looks reddish? Finally, in bringing in sophisticated counterfactual judgments about my own mind, the attractive idea that I can know that I see a hawk merely by attending to the hawk has been thrown overboard.

8.2. *Belief-Independence*

The known-illusion problem is entirely generated by the widespread assumption that, as Evans puts it, there is

a fundamental (almost defining) property of states of the informational system,³⁴ which I shall call their 'belief-independence': the subject's being in an informational state is independent of whether or not he believes that the state is veridical. It is a well-known fact about perceptual illusions that it will continue to

³² More generally, it should be *belief*, not just knowledge.

³³ For example, suppose one has a known-illusion of motion by viewing Kitaoka's "rotating snakes" figure (<http://visiome.neuroinf.jp/modules/xoonips/image.php?file_id=1173>). Assume, with Evans, that one does not believe that anything in the figure is moving. If one had not known anything about the figure beforehand, would one have judged that anything in the figure were moving? That depends. The figure and the motion both look so unusual, that a sensible person might well smell a rat. (Cf. Jackson 1977, 40-41.)

³⁴ Which subserves "perception, communication, and memory" and "constitutes the substratum of our cognitive lives" (1982, 122).

appear to us as though, say, one line is longer than the other (in the Müller-Lyer illusion), even though we are quite sure that it is not. (1982, 123)³⁵

Put in the present notation: even though one's visual system may (mis-)inform one that $[\dots x \dots]_v$, one may nonetheless resist its testimony and not believe this *v*-proposition.

But is it true that perception is belief-independent? Evans's correct observation about the Müller-Lyer illusion does not immediately establish this conclusion. He notes that it may appear to one that the lines are unequal even though one believes they are equal. For belief-independence to follow, it must also be assumed that if one believes that the lines are equal, one does not also believe that they are unequal. And since having contradictory beliefs is a familiar phenomenon, this assumption needs to be backed up with an argument.

Let us call the view that vision constitutively involves belief in the relevant *v*-proposition, *belief-dependence*. (Belief-dependence is, more-or-less, the "judgmental theory of perception" defended in Craig 1976.) Belief-dependence is not, it should be emphasized, the view that to enjoy visual appearances is simply to have beliefs of a certain sort. (For a reductive theory along these lines, see Armstrong 1961, ch.9.) Neither it is the view that perception can be analyzed partly in terms of belief. In these respects, belief-dependence is analogous to the view that knowledge constitutively involves belief: that does not imply that knowledge is belief, or that it can be analyzed partly in terms of belief. Although the passage from Evans does not conclusively establish that belief-dependence is false, it might be thought that the idea that one has contradictory beliefs in cases of known illusion is implausible. So can anything positive be said in favor of belief-dependence?

Here, briefly, are three considerations.³⁶ First, presumably some animals with visual systems very similar to ours (some other primates, say), cannot cognitively override visual illusions: in this sense, for them, seeing is always believing. Belief is thus built into their visual systems. And since we have basically the same visual systems, seeing is believing for us too.

Second, perception is clearly belief-like—which is why Armstrong-style attempts to reduce perception to belief were certainly worth trying. Perception compels belief; the visual appearance of unequal lines is accompanied by the belief that the lines are unequal, absent (apparent) evidence to the contrary. And perception has

³⁵ There is a slight infelicity in this passage. Evans has 'believes that the state is veridical' where it would have been better to write 'believes that *p*', where the proposition that *p* is (on his view) the "conceptualized" version of the content of the experience. (The vexed issue of conceptual versus nonconceptual content is not addressed here.)

³⁶ See also Byrne 2009, 450–451 and 2010, 120–121.

the same “direct of fit” as belief: false beliefs and illusory perceptions are mental states that are both failures, in some (admittedly obscure) sense. Belief-dependence explains both these features. Compulsion is explained simply because the visual appearance of unequal lines is always accompanied by the belief that the lines are unequal. Sometimes that belief will not be manifest because it is suppressed by the contrary belief that the lines are equal; remove that contrary belief, and one will have an unsuppressed belief that the lines are unequal, that will manifest itself in the usual way. And direction of fit is explained because the failure of a constitutive component of a perceptual state presumably implies the failure (or less than complete success) of the state as a whole.

Finally, consider the really quite remarkable phenomenon that numerous not-long-dead philosophers claimed to believe the deliverances of vision even in cases of illusion.³⁷ “When I see a tomato,” H. H. Price famously declared, “there is much I can doubt. I can doubt whether it is a tomato that I am seeing, and not a cleverly painted piece of wax. I can doubt whether there is any material thing there at all. Perhaps what I took for a tomato was really a reflection, perhaps I am even the victim of some hallucination. One thing however I cannot doubt; that there exists a red patch of a round and somewhat bulgy shape, standing out from a background of other color-patches, and having a certain visual depth” (Price 1932, 3). On the orthodox view, the plain man does not believe that there exists a bulgy red patch when he knows that the devious color-illusion has been explained. So why, on the orthodox view, do distinguished philosophers like Price believe the contrary after careful phenomenological study? Are they insane?

On some accounts of delusions (e.g., the Capgras delusion), they involve beliefs that are “modular” in something like the sense of Fodor 1983: delusory beliefs are largely inferentially isolated and persist despite evidence to the contrary (see, e.g., Jones, Delespaul, and van Os 2003). Belief-dependence offers a similar model of cases of known illusion. Since one believes the relevant *v*-proposition in a case of known illusion, one is in a position to (try to) follow SEE. Therefore, the known-illusion problem does not arise.³⁸

“modular”
beliefs

³⁷ More exactly, the deliverances of (roughly) mid-level vision. There are also many long-dead examples.

³⁸ If an ideal of rationality is avoidance of inconsistency, then belief-dependence implies that someone who suffers a perceptual illusion thereby falls short of the rational ideal. (As Craig 1976, 15–16 points out; see also Glüer 2009, 303 n. 10.) Can this be turned into a convincing objection?

No. It will not do simply to claim that the illuded subject is not, or need not be, irrational. Taken as a claim about a rational ideal, its truth is not evident. Taken as an ordinary sort of remark, on the other hand, it is true but not in conflict with belief-dependence. The belief that the subject knows to be false (e.g., a certain *v*-proposition that is true only if the lines are unequal) does not influence her verbal reports about the lengths of the lines, or any plans for action based on the

To close, three features of the transparency proposal should be highlighted. First, the transparency proposal is an *inferentialist* account of knowledge of what one sees; an odd inference, to be sure, but an inference nonetheless. Second, the transparency account is *economical*: it explains self-knowledge in terms of epistemic capacities and abilities that are needed for knowledge of other subject matters. And third, the account is *detectivist*: broadly causal mechanisms play an essential role in the acquisition of knowledge of what one sees, as they do generally in knowledge of contingent matters.

Economy is needed to explain why intelligent subjects with normal vision inevitably know what they see in the usual distinctive first-personal way.³⁹ Detectivism is just common sense. Inferentialism, on the other hand, runs counter to the usual characterization of self-knowledge as “direct” (that is, not inferential). Does this mean that the transparency proposal faces yet another serious objection? That is left as an exercise.⁴⁰

lengths of the lines. She is not therefore ‘irrational’ in the practical sense of an ordinary accusation of irrationality. The subject’s belief that the lines are unequal does little harm—at worst, it would make her a sense datum theorist. Indeed, given the epistemological account of this chapter, it actually does some good, by allowing the illuded subject to know what she sees.

Glüer, who thinks that belief-dependence founders on this sort of consideration, asserts that “there is nothing ‘irrational’ about the lines looking of different length” (2009, 303 n. 9). But she does not explain why this is true on the required reading of ‘irrational’.

³⁹ That is, it is needed to explain why dissociations of the sort mentioned at the end of section 6.1 do not occur.

⁴⁰ Many thanks to audiences at Stockholm University, Victoria University of Wellington, ANU, IJN, and NYU. I am especially grateful to Ned Block, Jesse Prinz, François Recanati, Nico Silins, Declan Smithies, Robert Stalnaker, and Daniel Stoljar.

REFERENCES

- Anscombe, G. E. M. 1965. "The Intentionality of Sensation: A Grammatical Feature." In *Analytical Philosophy*, second series, edited by R. J. Butler. Oxford: Blackwell.
- Armstrong, D. M. 1961. *Perception and the Physical World*. London: Routledge & Kegan Paul.
- . 1968. *A Materialist Theory of the Mind*. London: Routledge & Kegan Paul.
- Block, N. 1997. "Biology Versus Computation in the Study of Consciousness." *Behavioral and Brain Sciences* 20: 159–165.
- Brady, T. F., T. Konkle, G. A. Alvarez, and A. Oliva. 2008. "Visual Long-Term Memory Has a Massive Storage Capacity for Object Details." *Proceedings of the National Academy of Sciences* 105: 14325–14329.
- Burnham, R., and J. Clark. 1955. "A Test of Hue Memory." *Journal of Applied Psychology* 39: 164–172.
- Byrne, A. 2005. "Introspection." *Philosophical Topics* 33: 79–104.
- . 2009. "Experience and Content." *Philosophical Quarterly* 59: 429–451.
- . 2010. "Knowing That I am Thinking." In *Self-Knowledge*, edited by A. Hatzimoyisis. Oxford: Oxford University Press.
- . 2011. "Recollection, Perception, Imagination." *Philosophical Studies* 148: 15–26.
- Byrne, A., and D. R. Hilbert. 2003. "Color Realism and Color Science." *Behavioral and Brain Sciences* 26: 3–21.
- Byrne, A., and H. Logue. 2009. *Disjunctivism: Contemporary Readings*. Cambridge, Mass.: MIT Press.
- Chalmers, D. 1996. *The Conscious Mind*. Oxford: Oxford University Press.
- Craig, E. 1976. "Sensory Experience and the Foundations of Knowledge." *Synthese* 33: 1–24.
- Dretske, F. 1995. *Naturalizing the Mind*. Cambridge, Mass.: MIT Press.
- . 2003. "How Do You Know You Are Not a Zombie?" In *Privileged Access: Philosophical Accounts of Self-Knowledge*, edited by B. Gertler. Aldershot: Ashgate.
- Evans, G. 1982. *The Varieties of Reference*. Oxford: Oxford University Press.
- Fodor, J. A. 1983. *The Modularity of Mind*. Cambridge, Mass.: MIT Press.
- Gibson, J. J. 1979. *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.
- Glüer, K. 2009. "In Defence of a Doxastic Account of Experience." *Mind and Language* 24: 297–327.
- Goldenberg, G., W. Müllbacher, and A. Nowak. 1995. "Imagery Without Perception—A Case Study of Anosognosia for Cortical Blindness." *Neuropsychologia* 33: 1373–1382.
- Hartmann, J., W. Wolz, D. Roeltgen, and F. Loverso. 1991. "Denial of Visual Perception?" *Brain and Cognition* 16: 29–40.
- Hilbert, D. 1994. "Is Seeing Believing?" *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association* 1: 446–453.

- Jackson, F. 1977. *Perception: A Representative Theory*. Cambridge: Cambridge University Press.
- Jones, H., P. Delespaul, and J. van Os. 2003. "Jaspers Was Right After All—Delusions Are Distinct from Normal Beliefs." *British Journal of Psychiatry* 183: 285–286.
- Moore, G. E. 1903. "The Refutation of Idealism." *Mind* 7: 1–30.
- Peacocke, C. 1992. *A Study of Concepts*. Cambridge, Mass.: MIT Press.
- Price, H. H. 1932. *Perception*. London: Methuen.
- Recanati, F. 2007. *Perspectival Thought*. Oxford: Oxford University Press.
- Ryle, G. [1949] 1963. *The Concept of Mind*. London: Hutchinson. (Page reference to Peregrine Books 1963.)
- Searle, J. R. 1983. *Intentionality*. Cambridge: Cambridge University Press.
- Shoemaker, S. 1963. *Self-Knowledge and Self-Identity*. Ithaca, NY: Cornell University Press.
- . 1994. "Self-Knowledge and 'Inner Sense'." *Philosophy and Phenomenological Research* 54: 249–314.
- Sosa, E. 1999. "How to Defeat Opposition to Moore." *Philosophical Perspectives* 13: 141–153.
- Uchikawa, K., and M. Ikeda. 1986. "Accuracy of Memory for Brightness of Colored Lights Measured with Successive Comparison Method." *Journal of the Optical Society of America A* 3: 34–39.
- Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Wittgenstein, L. [1921] 2001. *Tractatus Logico-Philosophicus*. Translated by D. F. Pears and B. F. McGuinness. London: Routledge and Kegan Paul.